

## **Non-Fault Tolerant Network Nodes in a Multiple Fault Tolerant Network**

### **Field of the Invention**

The invention relates generally to computer networks, and more specifically to a method and apparatus providing non-fault tolerant network node operability in a fault-tolerant network.

Docket No. H16-26157

### **Notice of Copending Applications**

This application is related to the following copending applications, which are hereby incorporated by reference:

“Fault Tolerant Networking”, serial number 09/188,976; and  
Atty. docket number 256.044us1

### **Background of the Invention**

Computer networks have become increasingly important to communication and productivity in environments where computers are utilized for work. Electronic mail has in many situations replaced paper mail and faxes as a means of distribution of information, and the availability of vast amounts of information on the Internet has become an invaluable resource both for many work-related and personal tasks. The ability to exchange data over computer networks also enables sharing of computer resources such as printers in a work environment, and enables centralized network-based management of the networked computers.

For example, an office worker's personal computer may run software that is installed and updated automatically via a network, and that generates data that is printed to a networked printer shared by people in several different offices. The network may be used to inventory the software and hardware installed in each personal computer, greatly simplifying the task of inventory management. Also, the software and hardware configuration of each computer may be managed via the network, making the task of user support easier in a networked environment.

Networked computers also typically are connected to one or more network servers that provide data and resources to the networked computers. For example, a server may store a number of software applications that can be executed by the networked computers, or may store a database of data that can be accessed and utilized by the networked computers. The network servers typically also manage access to certain networked devices such as printers, which can be utilized by any of the networked computers. Also, a server may facilitate exchange of data such as e-mail or other similar services between the networked computers.

Connection from the local network to a larger network such as the Internet can provide greater ability to exchange data, such as by providing Internet e-mail access or access to the World Wide Web. These data connections make conducting business via the Internet practical, and have contributed to the growth in development and use of computer networks. Internet servers that provide data and serve functions such as e-commerce, streaming audio or video, e-mail, or provide other content rely on the operation of local networks as well as the Internet to provide a path between such data

servers and client computer systems.

But like other electronic systems, networks are subject to failures.

Misconfiguration, broken wires, failed electronic components, and a number of other factors can cause a computer network connection to fail, leading to possible inoperability of the computer network. Such failures can be minimized in critical networking environments such as process control, medical, or other critical applications by utilization of backup or redundant network components. One example is use of a second network connection to critical network nodes providing the same function as the first network connection. But, management of the network connections to facilitate operation in the event of a network failure can be a difficult task, and is itself subject to the ability of a network system or user to properly detect and compensate for the network fault. Furthermore, when both a primary and redundant network develop faults, exclusive use of either network will not provide full network operability.

One solution is use of a method or apparatus that can detect and manage the state of a network of computers utilizing redundant communication channels. Such a system incorporates in various embodiments nodes which are capable of detecting and managing the state of communication channels between the node and each other fault-tolerant network node to which it is connected. In some embodiments, such network nodes employ a network status data record indicating the state of each of a primary and redundant network connection to each other node, and further employ logic enabling determination of an operable data path to send and receive data between each

pair of nodes.

But, such networks will desirably include nodes which do not have full fault-tolerant capability. One common example of such a non-fault-tolerant network node is a standard office laser printer with a built-in network connection. What is needed is a method and apparatus to facilitate communication with non-fault-tolerant network nodes in such a fault-tolerant network system.

DRAFT DRAFT DRAFT DRAFT DRAFT DRAFT

### **Summary of the Invention**

The present invention provides a method and apparatus for operation of non-fault tolerant network nodes in a fault-tolerant network environment. In some embodiments, a network address or network location of any network nodes present on a network that are not fault-tolerant is determined and stored, and data to be sent to the detected non-fault-tolerant network nodes is routed only over that network to which the non-fault-tolerant network node is connected. In various further embodiments, the fault-tolerant network comprises a primary and redundant network with fault tolerant network nodes that are attached to each network; a non-fault-tolerant network node that is attached to either the primary or redundant network is then operable to communicate with any fault-tolerant network node via data sent over only the network to which the non-fault-tolerant network node is connected.

### **Brief Description of the Figures**

Figure 1 shows a diagram of a fault-tolerant computer network with multiple

fault-tolerant network nodes having primary and redundant network connections and having multiple non-fault-tolerant network nodes, consistent with an embodiment of the present invention.

Figure 2 shows a flowchart of a method of managing communication with non-fault-tolerant network nodes in a fault-tolerant computer network, consistent with an embodiment of the present invention.

PCT/US2012/045200

### **Detailed Description**

In the following detailed description of sample embodiments of the invention, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration specific sample embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized and that logical, mechanical, electrical, and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the invention is defined only by the appended claims.

The present invention provides a method and apparatus for managing communication with non-fault-tolerant network nodes in a fault-tolerant network. The invention is capable in various embodiments of identifying the network location and address of non-fault tolerant network nodes, and of routing data only over certain

portions of the network to communicate with the non-fault-tolerant network node. The network in some embodiments comprises a primary and a redundant network having connections to each fault-tolerant network node, and the invention comprises routing information to non-fault-tolerant network nodes connected to either the primary or redundant network via only that network to which the non-fault-tolerant network node is connected.

The invention in various forms is implemented within an existing network interface technology, such as Ethernet. In one such embodiment, the fault-tolerant network comprises two Ethernet connections connected to each fault-tolerant computer or node — a primary network connection and a redundant network connection. It is not critical for purposes of the invention which connection is the primary connection and which is the redundant connection, as the connections are physically and functionally similar. In the example embodiment discussed here, the primary and redundant network connections are interchangeable and are named primarily for the purpose of distinguishing the two networks from each other. Each of the primary and redundant networks also may have one or more non-fault-tolerant network nodes attached, and communication with such non-fault-tolerant networks is facilitated by the present invention.

Figure 1 illustrates an exemplary fault-tolerant network with fault-tolerant network nodes 101, 102 and 103. A primary network 104 and a redundant network 105 link each node to the other nodes of the network, as indicated by the lines connecting the nodes to each of the networks. Non-fault tolerant network nodes are

also connected to each network, including non-fault-tolerant network node 106 connected to the primary network 104 and non-fault-tolerant network node 107 connected to redundant network 105.

The fault-tolerant network connections linking the fault-tolerant network nodes are configured such that the fault-tolerant network nodes can communicate with each other despite multiple network faults, such as by use in some embodiments of particular node-to-node communication rules and network status monitoring capability within each node. The communication rules in various embodiments of a fault-tolerant network facilitate determination of a network path between each pair of nodes based on the network status data stored in and communicated between each fault-tolerant network node. Each fault-tolerant network node of such a system must be able to recognize non-fault-tolerant network nodes and adapt its communication rules for communicating with non-fault-tolerant network nodes such as nodes 106 and 107 of the example network of Figure 1.

In a typical single network configuration, data intended for a network node is simply sent over the network to the intended node. But, where multiple networks are combined to form a fault-tolerant network capable of compensating for multiple faults such as the network of Figure 1, data intended for a single non-fault-tolerant network node such as node 106 or 107 is desirably sent over the specific network to which the non-fault-tolerant network node is connected.

In some embodiments of the invention, data intended for a non-fault-tolerant network node such as non-fault-tolerant network node 106 is simply broadcast or

transmitted over both the primary network 104 and the redundant network 105, ensuring that the data is sent to the network to which the non-fault-tolerant network node is connected. Such a system does not require tracking addresses or locations of non-fault-tolerant network nodes, and simply relies on the network interface adapters of the redundant network 105 to filter out the extra data. But, such a configuration is reliant on the ability of the nodes connected to the redundant network 105 to ignore the data intended for a network node not attached to that network, and further wastes network bandwidth on the redundant network.

Other embodiments of the invention comprise maintaining an address table of detected non-fault-tolerant network nodes that are present on both the primary network 104 and the redundant network 105, and further associating each address or non-fault-tolerant network node with the network on which the node address was detected. In some embodiments of the invention, detection of the non-fault-tolerant network node address comprises monitoring for and intercepting Internet Protocol (IP) Address Resolution Protocol (ARP) packets that are sent by each node in certain IP-compatible network configurations. For example, each ARP packet in an Ethernet network contains the Media Access Control (MAC) address that uniquely identifies the node transmitting the IP ARP packet. The intercepted MAC address of each non-fault-tolerant network node is then recorded along with the network on which the non-fault-tolerant network node is detected. In other embodiments, other network hardware and communication protocols may be used for the same purpose, and are within the scope of the invention.

65  
64  
63  
62  
61  
60  
59  
58  
57  
56  
55  
54  
53  
52  
51  
50  
59

To send data from a fault-tolerant network node to a non-fault-tolerant network node in such embodiments of the invention, the address of the desired node is found in the stored address records of the sending fault-tolerant network node, and the associated network is determined. For example, if fault-tolerant network node 101 initiates a data transfer to non-fault-tolerant network node 106, node 101 searches its stored address records and finds the address of node 106, and further finds that the address data for node 106 was received on the primary network 104 rather than the redundant network 105. Node 101 then sends the data intended for node 106 only over network 104, eliminating the need to send the same data over redundant network 105 and use additional network bandwidth.

In further embodiments of the invention, fault-tolerant network nodes such as node 101 use the network status data indicating the ability of that node to communicate with other fault-tolerant network nodes to reroute data intended for a non-fault-tolerant network node around a network fault. This is achieved in some embodiments by initially sending the data on the network on which the non-fault-tolerant network node does not reside and using a selected fault-tolerant network node to transfer the sent data to the network on which the non-fault-tolerant network node resides at a point on the non-fault-tolerant network node's network such that the fault on the non-fault-tolerant network node's network is not between the transferring node and the non-fault-tolerant receiving node. Other embodiments exist in which data can be rerouted across the fault-tolerant networks to avoid multiple faults, and are within the scope of the invention.

In further embodiments, data sent to non-fault-tolerant network nodes is sent over all networks in the fault-tolerant network system rather than sent over a single network if the record containing address and network data for non-fault-tolerant network nodes does not contain data on the intended destination non-fault-tolerant network node. Sending such data comprises sending or replicating the data on both the primary and redundant network of the example network discussed above and shown in Figure 1.

Figure 2 is a flowchart of a method of managing communication between fault-tolerant network nodes and non-fault-tolerant network nodes in a fault-tolerant network such as the example network of Figure 1. At 201, each fault-tolerant network node determines the network address of any non-fault-tolerant network nodes present on each network to which the fault-tolerant network node is connected. This may be achieved in any suitable manner, including searching for IP ARP packets or other identifying data transmitted by the non-fault-tolerant network nodes. At 202, each fault-tolerant network node further determines the network on which each non-fault-tolerant network node exists. In some embodiments, this simply comprises detecting which network adapter in the detecting fault-tolerant network node detected the IP ARP packet or other identifying data. At 203, the fault-tolerant network nodes each store the data determined at 201 and 202. The address and network data for each non-fault-tolerant network node are associated with each other in the stored data in one embodiment, so that looking up a record for a particular non-fault-tolerant network node results in retrieval of both the network address of the node and the network on

which the node resides. In various embodiments, the process of determination of network addresses and networks associated with each non-fault-tolerant network node and the storing of this data is a continuous process, and occurs even during other operations such as execution of other blocks of the flowchart of Figure 2.

At 204, a fault-tolerant network node initiates sending data to a non-fault-tolerant network node. At 205, the stored data is searched for the address and network of the non-fault-tolerant network node. At 206, a decision is made based on determination of whether the address and network data for the non-fault-tolerant network node are present in the stored data. If the address and network data are present in the stored data, the data to be sent is sent from the fault-tolerant network node to the non-fault-tolerant network node over only that network to which the stored data indicates the non-fault-tolerant network node is connected at 207. In other embodiments, the data is sent indirectly via one or more intermediate nodes to the non-fault-tolerant network node, to avoid one or more network faults. If the address and network data are not present in the stored data, the data to be sent is sent over all networks to which the sending fault-tolerant network node is connected to ensure that the intended non-fault-tolerant network node receives the data. In the example of Figure 1, the data would be sent over both the primary network 104 and the redundant network 105.

The present invention provides a method and apparatus that enable a network with primary and redundant network connections to manage routing of data to non-fault-tolerant network nodes within the network. Some embodiments of the invention

incorporate a data record within each fault-tolerant network node that contains detected address and network data for each non-fault-tolerant network node, and which then is used by the fault-tolerant network node to determine over which network data intended for a specific non-fault-tolerant network node should be sent. In some embodiments, the invention includes rerouting data that cannot be transferred directly from a fault-tolerant network node to a non-fault-tolerant network node due to a network fault, and comprises routing the data to one or more intermediate nodes which are able to facilitate communication between the nodes.

Although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement which is calculated to achieve the same purpose may be substituted for the specific embodiments shown. This application is intended to cover any adaptations or variations of the invention. It is intended that this invention be limited only by the claims, and the full scope of equivalents thereof.